

The Xanadu storage system is a new form of software with potentially revolutionary implications.

This new form of storage has powerful new implications for personal computing, word processing, file management, the office of the future and its software, teleconferencing, electronic mail, electronic publishing, libraries of the future, and tomorrow's education.

People think the Xanadu system is complicated. Yet it is so much simpler than what's out there.

So-called "computer literacy" means learning to live, and being resigned to live, in a tangle of complexity that should never exist: of thousands of disconnected files and incompatible programs.

The Xanadu system maintains the connections between files (we call them "documents") and provides a platform of common data structure for all applications.

The Xanadu system is both a litmus and a Rorschach test. Everybody who has heard of it has an opinion on it, though perhaps five percent understand what it is. It seems that those who believe in it are either those who know nothing about computers and/or those who know *a lot*.

WHAT YOU HEAR: it's a plan for a worldwide network, intended to serve hundreds of millions of users simultaneously from the corpus of the world's stored writings, graphics and data. *This is true.*

THE IMPRESSION PEOPLE GET is of humungous hardware, centralized, requiring squillions of dollars to get off the ground. *This is wrong.*

The Xanadu system is the central software for serving users and running this network. It is a design for a new literature, and a *system of order* to make such a network understandable, usable, and readily expandible to any degree.

Most of all, it's new file software for keeping everything orderly and interconnected-- *in ways we need in our offices and personal computing lives.*

The program's prototype is operational now. No, not fast enough for millions yet. *But it's all there and will scale up cleanly.*

RATIONAL STORAGE, RADICAL STORAGE

Each byte is in principle stored only once (though for practical purposes there exist distributed copies throughout the network). And this permits our unique integration of links and commonalities.


By links we propose to map the true structure of data, and by commonalities we propose to map its interconnection. These explorable connections promise new understanding and control of the proliferating information that now swamps and confuses us.

This approach is not merely exotic. It is sensible. It keeps track of the origins, interconnections and true structure of everything.

It is "radical" in both senses of the term: dramatically different from what people are used to, and a return to the roots of what we want and what we do. We are so used to the structures and conventions of hierarchical files and limited compatibility that we do not see their inanity. What we call "computer literacy" and "computer basics" are really the willingness to live in this world of inane and unnecessary complication.

THE XANADU* PARADIGM

A Piece of Software that Proposes
A New Era of Computers,
A New Form of Instant Literature
And a Whole New World




1. COMMONALITY

COMMONALITY MANAGEMENT with the Xanadu System

In principle everything is stored *only once*.

A document is virtual, consisting of its *new* (or native) bytes and links, plus whatever previous bytes and links are included from other documents.




This extremely powerful mechanism, unique to the Xanadu system, has many desirable aspects, including:

CLARIFICATION OF ORIGINS (for text boilerplate, for management of business and scientific data).

VISIBILITY OF ALL LINKS: Any links to a given set of bytes are visible on *all* manifestations of these bytes, in whatever documents they may be included.

BRIDGES BETWEEN DOCUMENTS, across bytes in common, which may be followed by the user.




ALLOCATION OF CREDIT (in hypertext publishing)

ALLOCATION OF ROYALTY (in hypertext publishing).

2. LINKS


UNIVERSAL LINKAGE with the Xanadu System

The Xanadu link is an elemental unit maintained by the Xanadu software and addressed by its numerical ID, or found by various types of search.




The Xanadu link is a two-sided (or *dyadic*) connection between spans of bytes, documents and/or other links. (The one-sided or *monadic* link, a special case, is used for marking or identifying data.)

Two-sided links may be used as building blocks for more complex data structures. For instance, a comment comparing two text passages is constructed as a link to a link:



Links are specified to the Xanadu system by users or by application programs. A link, once specified, is maintained automatically by the Xanadu software, regardless of editorial changes which move, separate, or delete parts of its ends.

Any spans of bytes, in any document or file, on any server, may be linked to any other spans of bytes, in any document or file, on any server, by a link type which is unique or used elsewhere in the system.



All links are maintained by the system no matter how all documents and files change. THIS PROVIDES A STABLE METHOD FOR HANDLING LARGE COMPLEXES OF LINKED DATA FOR ALL PURPOSES.

3. LINKS ARE ON ALL COPIES

These two aspects taken together-- commonalities and links-- have extraordinary powers.

Most notably, this means that A LINK MADE TO ALL BYTES ANYWHERE IS PRESENT ON ALL MANIFESTATIONS OF THOSE BYTES, since there is in principle only one master copy of them.

For instance, suppose you read an office memorandum containing a certain piece of text that is native to another document-- say, a written description of a certain plan. Now you make a written comment on that plan. That comment is now available to the reader of the original, or to the reader of *any other document that includes that plan*.

Or if you place an electronic bookmark (a link) on a certain paragraph in version 1 of a document, and only one sentence of that paragraph remains in version 7, your bookmark will still be on it.

This also means that when the Xanadu system is used for widespread electronic publishing, people can publish links to other material which will be present in all different uses (by inclusion) of that material. This will be a powerfully unifying and clarifying aspect of the system.

4. CUMULATIVE ORDER

Everywhere else in the computer field, disorder seems to accumulate-- files with unknown contents and incomprehensible names, incompatibility, escalating complications when you try to tie things together.

But materials stored on the Xanadu system can be continually revised harmlessly in place, with all previous versions available (or ignorable). New versions and better links supersede the old, but you can always refer back to the way each part was before. This means that order and clarity can *increase*-- a unique capability.

5. A UNIVERSAL DATA STRUCTURE

By using links to mark and type data elements, and to represent typed connections between the data elements, the Xanadu system provides A UNIVERSAL DATA STRUCTURE TO WHICH ALL OTHER DATA MAY BE MAPPED, allowing coordinated access to a variety of application programs whose development need not be coordinated.

(Note that because material is left in place once it is stored, the Xanadu software is particularly well adapted to write-once and read-only memory-- i.e., WORM drives and CD-ROM.)

READ ALL ABOUT IT

Theodor H. Nelson, *Literary Machines* edition 87.1. \$25 postpaid from Project Xanadu, 8480 Fredericksburg #138, San Antonio TX 78229 (add \$5 for purchase order, \$5 for overseas airmail).

Theodor H. Nelson, "Replacing the Printed Word." In Proceedings of the 1980 IFIP World Computer Conference.

Theodor H. Nelson, "All for One and One for All." Proceedings of Hypertext 87 Conference.

Theodor H. Nelson, "The Tyranny of the File." *Datamation*, 15 Dec 86.

Theodor H. Nelson, "Immense Storage Management." *Byte*, January 1988.

Theodor H. Nelson, *Computer Lib*, second edition. Tempus Books (Microsoft Press), 1987.

Howard Rheingold, *Tools for Thought*. Simon & Schuster, 1985.

K. Eric Drexler, *Engines of Creation*. Doubleday, 1986.

URGENT CORRECTION

(Note: OMNI magazine, and others apparently using it as a source, recently said that only the original author of a Xanadu document will be able to make links to or from it. This is CATEGORICALLY FALSE. In the Xanadu publishing method, anyone will be able to publish links to or from any already-published material.)


*"Xanadu," the Eternal-Flaming-X and "Silverstand" are trade marks for hypertext and software products and designs offered by Project Xanadu, and service marks for storage, retrieval, data addressing, transmission, transaction and publishing methods, consulting services and new paradigms offered by Project Xanadu.

However, some people feel more comfortable calling this idea by its generic term, *xanalogical storage*.

Xanalogical storage consists of pooled material whose elements may be simultaneously used in any number of new units, and where links may be made between arbitrary portions of any units.

We know of no other system for doing this at the present time.

XANALOGICAL STORAGE



THE XANADU PUBLISHING METHOD

Because material may be freely included from any document to any other, the Xanadu system provides a simple and elegant solution to various problems of electronic publishing, under a simple set of agreements for participation.

COPYRIGHT is to be handled simply within the system. Royalty will be paid automatically by the reader, or anyone making a printout, on a byte-by-byte basis: each bytesworth of royalty (or *nib*) is paid to the owner of its native document. Credit for originating material is easily allocated, since all quotations may be followed "sideways" to their documents of origin.

No permission is necessary to include any part of a document in another, since inclusions are virtual and require no copying.

No one need participate in this system, or give permission for the use of their documents in Xanadu publication; but we anticipate that the power, cleanliness and universality of these arrangements will gradually gather momentum throughout the world.

We are proposing a clean, simple, universal and fair system of electronic literature that builds on the structures and relationships (quotation and citation, credit and royalty) already present in literature as we know it. There is a place for author, editor, publisher, journal and so on, transposing their respective roles to this new world. And it will make us all clearer-minded and smarter.

This has striking implications. The power of being able to follow links on a great web, to "own" a whole universe of literature (by one's private bookmarks), and freely to modify what one reads into new versions of one's own-- all these promise new and unprecedented freedoms for the student and scholar, businessman and scientist, and indeed for anyone who works with information.

WON'T WE DROWN IN LINKS?

You only see what you're interested in. Links may be filtered by type and origin as the user desires.

CONTENT CONTROL VERSUS CONTEXT CONTROL

The owner of a document maintains *content control*: no one else may change the document or its daughter versions without the owner's permission.

However, others may link to it and quote parts of it, meaning a relinquishment of *context control*. But a moment's reflection will make it clear that it cannot be otherwise. In the era of electronic publishing, what is called to the user's screen may be placed in any context of other windows and materials on that screen; moreover, the user may take the bytes of published materials and put them to other use. This cannot be prevented and must be considered legitimate.

What the Xanadu system does, however, is to *keep this use connected*. Thus we can keep track of the origin and original context of all materials, and any links to them, no matter what other contexts they go into; and always view the originals side-by-side. This is, we believe, the best possible solution for all parties.

LOCAL NETWORKING WITH THE XANADU SYSTEM

The Xanadu storage module should find ready application as a storage server for local-area networks. Files (documents) may intrinsically be distributed across servers.

APPLICATION DESIGN FOR THE XANADU ENGINE

This changes all the rules.

Xanadu-compatible software need do no storage management, except as data-modification calls to and from the system. The application designer is freed of a great number of the usual concerns; and thus may turn to concentrate on the true virtuality of the particular application: issues of link design, interaction, appearance and conceptual structure.

The programmer of front-end applications need deal with no system calls other than those to the Xanadu module (the FEBE, or Front End-Back End protocol). The Xanadu module in turn handles device communication to storage units, storage allocation and addressing, and communication between Xanadu servers.

Note that even if the back-end code is bound into an application or front-end program, communication with it is still through the FEBE protocol.

A "front end" to the Xanadu system is any program that uses it.

While we have done some modest front-end work, the Xanadu team is not really in the front-end business. We prefer to leave to this to our protocol partners.

There is no Right Way for screens to look or act with Xanadu data. Indeed, screens will be more opulent every year, and there will be many new forms of interaction-- probably making today's personal computers seem barbaric. So every experiment and approach must be considered valid, but no front-end style may be considered final.

USE OF THE XANADU STORAGE MODULE IN YOUR SOFTWARE

The Xanadu storage module is in C and typically compiles to about 135K in 68000 code; compilation figures will naturally vary for different compilers and different target machines. It wants as many megabytes of workspace as it can get. The protocol manager ("the back of the front") compiles to about 30K and wants about a megabyte of work space.

We are looking for collaborators and clients in two categories: PROTOCOL PARTNERS, who want to work with the Xanadu system at the protocol level, and need no non-disclosure agreement; and ENFILADE PARTNERS, who wish to bind copies of the Xanadu module into application programs or otherwise use aspects of enfilaide technology. For this a non-disclosure agreement is required.

WELCOME TO OUR ADDRESS SPACE

Xanadu storage offers a way to bring compatibility to the Babel of the computer world.

And even to the new Babels. Today's new enthusiasm for hypertext and hypermedia is creating a slew of newly incompatible hypertext systems. Each designer thinks she or he has found the best way to handle link types, or interaction, or whatever. Mapping these to Xanadu storage should make it possible for all types of hyperdocument to coexist and interlink.

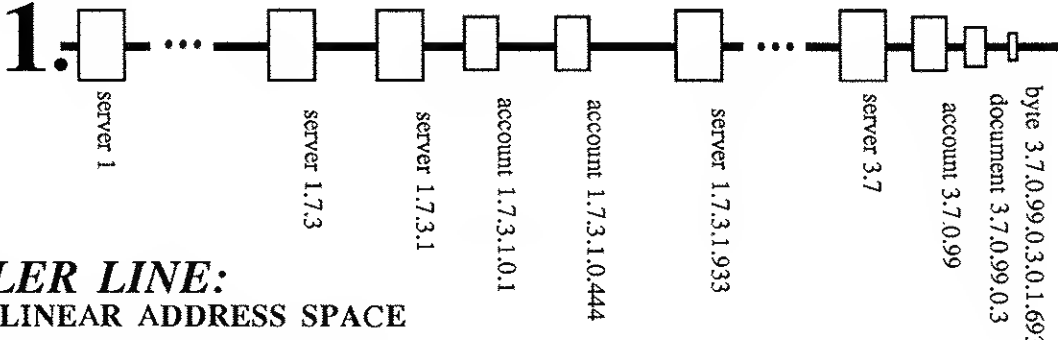
All servers of unchanging data can easily be assigned positions in the Xanadu address space, making it possible to embrace their materials logically without containing them physically. Servers offering cooperative connection can have their materials delivered through Xanadu servers, and vice versa.

The Xanadu address structure (invented by Mark S. Miller and further developed for this purpose by Miller with Roger Gregory) consists of ADDRESSES and SPANS, specified by an unusual new kind of number we call a *tumbler*. Tumblers are the principal currency of the FEBE protocol.

A tumbler consists of four fields, designating Server or Node, User or Account, Document/Version, and Element (byte or link). Each of these fields consists of as many integers as needed, separated by points. The fields themselves are separated by a zero between points, which has arithmetic significance.

Tumbler addresses may be arrayed in a line; this is called, of course, *the tumbler line*. Unlike the integer line or real line, it depends on what addresses have been assigned, and consists of all those tumbler addresses which are so far taken.

prefix permits addressing whole universe at once




THE TUMBLER LINE: A WRITE-ONCE LINEAR ADDRESS SPACE

Linear sequence permits whole spans to be referenced.

This linear array in turn permits the specification of *tumbler spans*, which include all existing addresses between the first element of the span and the first element after the span.

The rules for tumbler subdivision, addition and subtraction (permitting the specification of spans) are explained in *Literary Machines* 87.1 and "Immense Storage Management."

We may draw a square of tumbler lines, providing a way to map and visualize the parameters of links and of search requests to the Xanadu system.



LINKS ON THE TUMBLER SQUARE. The tumbler square maps all four parameters of a link: its left or from-spans, its right or to-spans, type or three-spans, and its home, which is a single element in a given document.

SEARCH REQUESTS ON THE TUMBLER SQUARE. A principal type of Xanadu search looks for all the links with specific parameters. Such a search may also be mapped to the tumbler square, and acts as a sieve to find all the eligible links. To be found, a link must satisfy one user-specified span OR another on each side of the square; by satisfying all sides, it satisfies the AND of the ORs.

YOU KNOW WHERE TO FIND US

Contact for general information and consulting:

Theodor H. Nelson, Director
Project Xanadu
8480 Fredericksburg, Suite 138
San Antonio TX 78229

Phone: 512/ 692-7346

Contact for technical details: (beyond what's here) and OEM terms:

Roger Gregory, President;
XOC, Inc., Box 7213
Menlo Park CA 94026.

Phone: 408/ 244-2643